



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

| APPLICATION NO. | FILING DATE | FIRST NAMED INVENTOR | ATTORNEY DOCKET NO. | CONFIRMATION NO. |
|--|-------------|------------------------|----------------------|------------------|
| 10/676,478 | 09/30/2003 | Ali-Reza Adl-Tabatabai | 42P17411 | 7528 |
| 8791 | 7590 | 10/10/2007 | EXAMINER | |
| BLAKELY SOKOLOFF TAYLOR & ZAFMAN 1279 OAKMEAD PARKWAY SUNNYVALE, CA 94085-4040 | | | KROFCHECK, MICHAEL C | |
| | | ART UNIT | PAPER NUMBER | |
| | | 2186 | | |
| | | MAIL DATE | DELIVERY MODE | |
| | | 10/10/2007 | PAPER | |

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

| | | | |
|------------------------------|------------------------|----------------------|--|
| Office Action Summary | Application No. | Applicant(s) | |
| | 10/676,478 | ADL-TABATABAI ET AL. | |
| | Examiner | Art Unit | |
| | Michael Krofcheck | 2186 | |

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --
Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) Responsive to communication(s) filed on 29 August 2007.
- 2a) This action is FINAL. 2b) This action is non-final.
- 3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) Claim(s) 1-33 is/are pending in the application.
 - 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) Claim(s) _____ is/are allowed.
- 6) Claim(s) 1-33 is/are rejected.
- 7) Claim(s) _____ is/are objected to.
- 8) Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) The specification is objected to by the Examiner.
- 10) The drawing(s) filed on 30 September 2003 is/are: a) accepted or b) objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
 - a) All b) Some * c) None of:
 1. Certified copies of the priority documents have been received.
 2. Certified copies of the priority documents have been received in Application No. _____.
 3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|--|---|
| 1) <input checked="" type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413) |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | Paper No(s)/Mail Date: _____ |
| 3) <input type="checkbox"/> Information Disclosure Statement(s) (PTO/SB/08) | 5) <input type="checkbox"/> Notice of Informal Patent Application |
| Paper No(s)/Mail Date _____ | 6) <input type="checkbox"/> Other: _____ |

DETAILED ACTION

1. This office action is in response to the RCE filed on 8/29/2007.
2. Claims 1, 15, 24, and 31 have been amended.
3. The objections/rejections from the prior correspondence not restated herein have been withdrawn.

Claim Rejections - 35 USC § 103

4. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

5. The factual inquiries set forth in *Graham v. John Deere Co.*, 383 U.S. 1, 148 USPQ 459 (1966), that are applied for establishing a background for determining obviousness under 35 U.S.C. 103(a) are summarized as follows:

1. Determining the scope and contents of the prior art.
2. Ascertaining the differences between the prior art and the claims at issue.
3. Resolving the level of ordinary skill in the pertinent art.
4. Considering objective evidence present in the application indicating obviousness or nonobviousness.

6. Claims 1-3, 8, 10, 14-17, 21, 23-25, and 31-33 are rejected under 35 U.S.C. 103(a) as being unpatentable over Shimoi et al., U.S. Patent 5,652,857 (hereinafter

Shimoi), Corcoran et al., U.S. Patent 6,449,689 (hereinafter Corcoran), and Kim et al., US patent 6859870.

7. With respect to claim 1, Shimoi teaches of a computer system comprising: a central processing unit (CPU) (fig. 1; item 16; column 8, lines 3 - 9); and

a cache memory, coupled to the CPU (fig. 1; item 28; column 8, lines 3 - 9), including: a main cache having a plurality of cache lines, each of the plurality of cache lines being compressible to form compressed cache lines to store additional data (fig. 4, 5, 7; column 10, lines 20 – 22; column 11, lines 2 – 11; where the cache memory (main cache) is divided into a non-compression data area and a compression data area. The logic block numbers (cache lines) #a, #b, #c, and #d are each compressed (not compressed cache lines) into a compression group. The compressed cache contains cache blocks 70-1 – 70-n (a compressed cache lines of cache lines));

a cache controller having compression logic to determine that a cache line is to be combined with a resident companion cache line to form the compressed cache line if the companion cache line is resident in the cache memory and to store the compressed cache line in the cache line of the resident companion (fig. 4, 7, items 48, 50, 52; column 10, lines 7-15, line 61-column 11, line 12; the compression circuit receives each block (cache line) from the cache memory and compresses each block and stores them as a compression group (compressed cache line). As the individual blocks are located adjacent to one another, shown in fig. 7, and are compressed into the same group, they are “companion” lines. The compression group is stored in the cache line of the

resident companion cache line as the originally uncompressed cache lines are all apart of the compression group).

Shimoi fails to specifically teach of a plurality of storage pools to hold a segment of the additional data for a compressed cache line.

However, Corcoran teaches of compressing data lines in a memory (fig. 2; column 4, lines 44 - 59) and a plurality of storage pools to hold a segment of the additional data for a compressed line (fig. 2; items 25; column 4, lines 44 – 49, and 60 – column 5, line 2; where any data blocks that fail to compress at least average will be partially stored in the overflow partition (storage pool)).

Kim teaches of a cache line retrieved from a main memory device is to be compressed in a cache (column 1, lines 55-67, column 3, lines 41-54).

It would have been obvious to one of ordinary skill in the art having the teachings of Shimoi and Corcoran at the time of the invention to incorporate the structure, and methods of operation of a dual partitioned storage device, the first partition comprising a compressed data track/line area and the second partition comprising an overflow area corresponding to those compressed data tracks, from Corcoran into the already partitioned and compressed cache memory of Shimoi. The motivation for this would have been to provide a more efficient data compression system that does not waste storage space (Corcoran column 1, lines 33 – 41; column 2, lines 18-19).

It would have been obvious to one of ordinary skill in the art having the teachings of Shimoi, Corcoran, and Kim at the time of the invention to also read the cache lines from main memory before compressing them into the cache of the combination of

Shimoi and Corcoran, since it is known in the art to retrieve information from the main memory into the cache memory for the processor to access.

8. With respect to claim 15, Shimoi teaches of a cache memory comprising: a main cache having a plurality of cache lines, each of the plurality of cache lines being compressible to form compressed cache lines to store additional data (fig. 4, 5, 7; column 10, lines 20 – 22; column 11, lines 2 – 11; where the cache memory (main cache) is divided into a non-compression data area and a compression data area. The logic block numbers (cache lines) #a, #b, #c, and #d are each compressed (not compressed cache lines) into a compression group. The compressed cache contains cache blocks 70-1 – 70-n (a compressed cache lines of cache lines)).

wherein a cache line is determined to be compressible if a companion cache line of a cache line is resident in the cache memory, and wherein the retrieved cache line and the companion cache line are combined and stored in the cache line of the resident companion if the cache line is determined to be compressible (fig. 4, 7; column 10, lines 7-15, line 61-column 11, line 12; the compression circuit receives each block (cache line) from the cache memory and compresses each block and stores them as a compression group (compressed cache line). The compression group is stored in the cache line of the resident companion cache line as the originally uncompressed cache lines are all apart of the compression group);

Shimoi fails to specifically teach of a plurality of storage pools to hold a segment of the additional data for a compressed cache line.

However, Corcoran teaches of compressing data lines in a memory (fig. 2; column 4, lines 44 - 59) and a plurality of storage pools to hold a segment of the additional data for a compressed line (fig. 2; items 25; column 4, lines 44 – 49, and 60 – column 5, line 2; where any data blocks that fail to compress at least average will be partially stored in the overflow partition (storage pool)).

Kim teaches of a cache line retrieved from a main memory device is to be compressed in a cache (column 1, lines 55-67, column 3, lines 41-54).

9. With respect to claim 24, Shimoi teaches of a method comprising: compressing one or more of a plurality of cache lines to form one or more compressed cache lines to store additional data by: determining if a companion cache line of a cache line is resident in a cache memory, and wherein the retrieved cache line and the companion cache line are combined (fig. 4, 5, 7; column 10, lines 20 – 22; column 10, line 61 – column 11, line 12; the logic block numbers (cache lines) #a, #b, #c, and #d are each compressed into a compression group. The cache blocks 68-1 – 68-4 are compressed into compression group 70-n);

combining the retrieved cache line with the companion cache line if the companion cache line is determined to be compressible (fig. 4, 7; column 10, lines 7-15, line 61-column 11, line 12; the compression circuit receives each block (cache line) from the cache memory and compresses each block and stores them as a compression group (compressed cache line));

storing the combined cache line in the cache line of the resident companion (fig. 4, 7; column 10, lines 7-15, line 61-column 11, line 12; The compression group is stored

in the cache line of the resident companion cache line as the originally uncompressed cache lines are all apart of the compression group).

Shimoi fails to specifically teach of storing a second component of the data in one or more of a plurality of storage pools.

However, Corcoran teaches of storing a first component of the data in a compressed memory (figs. 3A-C; column 5, lines 41 – 48; where the first portion of the compressed data is stored in a slot in the storage partition)

storing a second component of the data in one or more of a plurality of storage pools (figs. 3A-C; column 5, lines 41 – 48; where the second portion of the compressed data is stored in a slot in the overflow partition (storage pools)).

Kim teaches of a cache line retrieved from a main memory device is to be compressed in a cache (column 1, lines 55-67, column 3, lines 41-54).

10. With respect to claim 31, Shimoi teaches of a computer system comprising: a central processing unit (CPU) (fig. 1; item 16; column 8, lines 3 - 9); and

a cache memory, coupled to the CPU (fig. 1; item 28; column 8, lines 3 - 9), including: a main cache having a plurality of cache lines, each of the plurality of cache lines being compressible to form compressed cache lines to store additional data (fig. 4, 5, 7; column 10, lines 20 – 22; column 11, lines 2 – 11; where the cache memory (main cache) is divided into a non-compression data area and a compression data area. The logic block numbers (cache lines) #a, #b, #c, and #d are each compressed (not compressed cache lines) into a compression group. The compressed cache contains cache blocks 70-1 – 70-n (a compressed cache lines of cache lines)), and

a main memory device coupled to the CPU (fig. 1; item 20; column 8, lines 3 - 9)

a cache controller having compression logic to determine that a cache line is to be combined with a resident companion cache line to form the compressed cache line if the companion cache line is resident in the cache memory and to store the compressed cache line in the cache line of the resident companion (fig. 4, 7, items 48, 50, 52; column 10, lines 7-15, line 61-column 11, line 12; the compression circuit receives each block (cache line) and compresses each block and stores them as a compression group (compressed cache line). The compression group is stored in the cache line of the resident companion cache line as the originally uncompressed cache lines are all apart of the compression group).

Shimoi fails to specifically teach of a plurality of storage pools to hold a segment of the additional data for a compressed cache line.

However, Corcoran teaches of compressing data lines in a memory (fig. 2; column 4, lines 44 - 59) and a plurality of storage pools to hold a segment of the additional data for a compressed line (fig. 2; items 25; column 4, lines 44 – 49, and 60 – column 5, line 2; where any data blocks that fail to compress at least average will be partially stored in the overflow partition (storage pool)).

Kim teaches of a cache line retrieved from a main memory device that is coupled to the CPU is to be compressed in a cache (column 1, lines 55-67, column 3, lines 41-54).

11. With respect to claims 2, 16, and 32, Shimoi fails to specifically teach of wherein each of the plurality of storage pools include a plurality of fixed width storage fields.

However, Corcoran teaches of wherein each of the plurality of storage pools include a plurality of fixed width storage fields (fig. 2; column 4, lines 44 – 49; where the blocks are a fixed-size).

12. With respect to claims 3, 17, 25, and 33, Shimoi teaches of a plurality of cache lines (fig. 7, items 70-1 – 70-n). Shimoi fails to explicitly teach of wherein the plurality of cache lines are included within a plurality of sets.

However, Corcoran teaches of wherein the plurality of data lines are included within a plurality of sets (fig. 2, column 4, lines 44 – 49; where one or more tracks (sets) contain data blocks (lines) B1-Bx).

13. With respect to claims 8 and 21, Shimoi fails to explicitly teach of wherein a storage pool is shared by two or more of the plurality of sets.

However, Corcoran teaches of wherein a storage pool is shared by two or more of the plurality of sets (fig. 1; column 5, lines 62 – 66; in response to an overflow situation, the controller finds empty slots in the overflow partition and provides the address of the slots where overflow data can be written. Therefore, the slots in the overflow partition are shared by all of the tracks (sets)).

14. With respect to claims 14 and 23, Shimoi fails to explicitly teach of wherein a storage pool is shared by all of the plurality of sets.

However, Corcoran teaches of wherein a storage pool is shared by all of the plurality of sets (fig. 1; column 5, lines 62 – 66; in response to an overflow situation, the controller finds empty slots in the overflow partition and provides the address of the

slots where overflow data can be written. Therefore, the slots in the overflow partition are shared by all of the tracks (sets)).

15. With respect to claim 10, Shimoi teaches of a cache controller coupled to the cache memory (fig. 1; item 26; column 8, lines 3 - 9).

16. Claims 4 - 7, 9, 18 - 20, 22, and 26 - 30 rejected under 35 U.S.C. 103(a) as being unpatentable over Shimoi, Corcoran and Kim, as applied to claims 3, 8, 17, 21, and 25, respectively, and further in view of Obara U.S. Patent 6,115,787 (hereinafter Obara).

17. With respect to claims 4, 18, and 26, The combination of Shimoi, Corcoran and Kim fail to specifically teach of wherein a storage pool is allocated to each of the plurality of sets.

However, Obara teaches of wherein a storage pool is allocated to each of the plurality of sets (fig. 1; column 10, lines 41 – 59; where the second of each cache segment pair (storage pool) is used to store overflow data unable to be stored in the primary cache segment (set)).

It would have been obvious to one of ordinary skill in the art having the teachings of Shimoi, Corcoran, Kim, and Obara at the time of the invention to pair the overflow blocks in the combination of Shimoi, Corcoran and Kim with individual tracks the way an overflow cache segment is paired with a primary cache segment in Obara. The motivation for this would have been to more efficiently operate the cache memory (Obara column 1, lines 10 – 26) by guaranteeing that a cache segment has space for an overflow if needed.

18. With respect to claims 6 and 20, The combination of Shimoi Corcoran, and Kim fail to specifically teach of wherein multiple storage fields within each storage pool is allocated a line within one of the plurality of sets.

However, Obara teaches of wherein multiple storage fields within each storage pool is allocated a line within one of the plurality of sets (fig. 1; column 12, lines 11-43; where each block in the secondary cache segment (storage field) corresponds to a cache block (cache line) in the primary cache segment (set)).

It would have been obvious to one of ordinary skill in the art having the teachings of Shimoi, Corcoran, Kim, and Obara at the time of the invention to link the overflow blocks to the cache storage blocks in the combination of Shimoi, Corcoran, and Kim as is done in Obara. The motivation for this would have been to eliminate the need to store pointers in the cache memory to indicate the corresponding block, which takes up valuable space and overhead (Obara column 12, lines 29 - 43).

19. With respect to claim 7, The combination of Shimoi, Corcoran, and Kim fail to specifically teach of wherein each storage field mapped to one of the plurality of sets is sorted according to a logical ordering.

However, Obara teaches of wherein each storage field mapped to one of the plurality of sets is sorted according to a logical ordering (fig. 1; column 12, lines 11-43; where each block in the secondary cache segment (storage field) is ordered to match the blocks (cache lines) in the primary cache segment (sets)).

20. With respect to claims 5, 19, and 27, Corcoran teaches of wherein an indicator is associated with each storage field of a storage pool to indicate a line within one of the

plurality of sets to which a storage field is assigned (fig. 2; column 4, lines 50 – 59; where the address pointers (indicators) are “associated” with the overflow slots (storage fields). By being stored in their data block (line) the address pointers “indicate” the data block within the track (set) that corresponds to the overflow slots).

21. With respect to claim 28, The combination of Shimoi, Corcoran, and Kim fail to specifically teach of allocating a storage pool to a line within one of the plurality of sets.

However, Obara teaches of allocating a storage pool to a line within one of the plurality of sets (fig. 1; column 12, lines 11-43; where each block in the secondary cache segment (storage pool) corresponds to a cache block (cache line) in the primary cache segment (set)).

22. With respect to claim 29, The combination of Shimoi, Corcoran, and Kim fail to specifically teach of mapping each storage field to one of the plurality of sets.

However, Obara teaches of mapping each storage field to one of the plurality of sets (fig. 1; column 12, lines 11-43; where each block in the secondary cache segment (storage field) corresponds to a cache block (cache line) in the primary cache segment (set). Each cache block in the secondary cache segment (storage field) therefore is mapped to that primary cache segment (set)).

23. With respect to claims 9, and 22, The combination of Shimoi, Corcoran, and Kim fails to specifically teach of wherein an indicator is associated with each line of a storage pool to indicate which of the plurality of sets to which a storage field is assigned.

However, Obara teaches of wherein an indicator is associated with each line of a storage pool to indicate which of the plurality of sets to which a storage field is assigned

(figs. 1-4; column 11, lines 6 – 48; where the primary SGCB entry indicates the address for the primary cache segment (set) as well as the address for the secondary SGCB entry which in turn contains the address for the secondary cache segment (storage pool). The blocks (lines) within the secondary cache segment correspond to the same blocks (lines) within the primary cache segment; thereby being linked through the SGCB entries).

It would have been obvious to one of ordinary skill in the art having the teachings of Shimoi, Corcoran, Kim, and Obara at the time of the invention to link the overflow blocks to the cache storage blocks in the combination of Shimoi, Corcoran and Kim by the SGCB entries as is done in Obara. The motivation for this would have been to eliminate the need to store pointers in the cache memory to indicate the corresponding block, which takes up valuable space and overhead (Obara column 12, lines 29 - 43).

24. With respect to claim 30, The combination of Shimoi, Corcoran, and Kim fails to specifically teach of associating an indicator with each line of a storage pool to indicate which of the plurality of sets to which a storage field is assigned.

However, Obara teaches of associating an indicator with each line of a storage pool to indicate which of the plurality of sets to which a storage field is assigned (figs. 1-4; column 11, lines 6 – 48; where the primary SGCB entry indicates the address for the primary cache segment (set) as well as the address for the secondary SGCB entry which in turn contains the address for the secondary cache segment (storage pool). The blocks (lines) within the secondary cache segment correspond to the same blocks

(lines) within the primary cache segment; thereby being linked through the SGCB entries).

25. Claims 11-13 are rejected under 35 U.S.C. 103(a) as being unpatentable over Shimoi, Corcoran and Kim, as applied to claim 10 above, and further in view of Cypher, U.S. Patent 6,629,205 (hereinafter Cypher).

26. With respect to claim 11, Shimoi teaches of a cache controller to accesses the cache (fig. 1; column 8, lines 3 - 9). Shimoi fails to explicitly teach of accessing cache lines and storage pools in parallel.

However, Cypher teaches of accessing different cache memory partitions in parallel (fig. 1; column 4, lines 39 – 50; column 7, lines 42 – 46; where the cache controller is capable of simultaneously accessing cache tags associated with different cache classes).

It would have been obvious to one of ordinary skill in the art having the teachings of Shimoi, Corcoran, Kim, and Cypher at the time of the invention to apply the concept and means for simultaneous accessing multiple cache memory partitions as taught in Cypher in the cache controller of the combination of Shimoi, Corcoran. The motivation for this would have been to increase the snoop bandwidth, thus increasing the efficiency of the cache memory (Cypher column 2, lines 49–58).

27. With respect to claim 12, the combination of Shimoi, Corcoran, Kim, and Cypher teach of all the limitations of the parent claims as discussed supra. The combination of Shimoi, Corcoran, fails to explicitly teach of wherein accessing the cache lines and

storage pools in parallel comprises the cache controller simultaneously dispatching set bits to the cache lines and storage pools.

Cypher also teaches of wherein accessing different cache memory partitions in parallel comprises simultaneously dispatching set bits to the different cache memory partitions (figs. 4, 5; column 6, lines 13 – 23; where the controller concurrently conveys the index addresses (set bits) to the separate memory partitions).

28. With respect to claim 13, Shimoi fails to specifically teach of wherein the cache controller accesses the cache lines and storage pools serially.

However, Corcoran teaches of wherein the controller accesses the data lines and storage pools serially (figs. 1, 3a-c; column 5, lines 41 – 65; where the data is compressed and too large to fit in only the slot in the storage partition. The 1st part of the compressed data is stored in the storage partition and then the disk controller locates an empty slot in the overflow partition and provides the address of the empty slots. The remaining data is then stored in the overflow slots).

Response to Arguments

29. Applicant's arguments with respect to the claims have been considered but are moot in view of the new ground(s) of rejection.

Conclusion

30. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Michael Kroccheck whose telephone number is 571-272-8193. The examiner can normally be reached on Monday - Friday.
31. If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Matt Kim can be reached on 571-272-4182. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.
32. Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).



Michael Kroccheck



MATTHEW KIM
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100